

Express Mailing Label No. EV 406340122 US

PATENT APPLICATION
Docket No. 3408.2.4

UNITED STATES PATENT APPLICATION

of

DAVID A. EATOUGH

for

CLIENT INITIATED MULTICAST DOMAIN DISCOVERY

CLIENT INITIATED MULTICAST DOMAIN DISCOVERY

TECHNICAL FIELD

[01] The present invention relates generally to computer system administration. More specifically, the present invention relates to client initiated multicast domain discovery.

BACKGROUND

[02] Computer and communication technologies continue to advance at a rapid pace. Indeed, computer and communication technologies are involved in many aspects of a person's day. For example, many devices being used today have a small computer inside of the device. These small computers come in varying sizes and degrees of sophistication. Computers commonly used include everything from hand-held computing devices to large multi-processor computer systems.

[03] Computers are used in almost all aspects of business, industry and academic endeavors. More and more homes are using computers as well. The pervasiveness of computers has been accelerated by the increased use of computer networks, including the Internet. Most companies have one or more computer networks and also make extensive use of the Internet. The productivity of employees often requires human and computer interaction. Improvements in computers and software have been a force for bringing about great increases in business and industrial productivity.

[04] Maintaining and supporting computer systems is important to anyone who relies on computers. Whether a computer or computing device is in a home or at a business, at least some maintenance and/or support is often needed. For example, sometimes there are problems with computer hardware. In addition, computer hardware is often upgraded and replaced with new components. Similarly computer software is also frequently upgraded or replaced. New computer hardware and software is continually being integrated into systems across the world.

[05] Installing new computer hardware and/or software, or problems with existing systems, may cause down-time during which the business or individual operate at diminished levels or are

otherwise unable to use the computers as they typically do. Most individuals and businesses try to minimize computer problems so as to avoid down-time.

[06] When a business or individual is trying to decide whether to make a change to a computer system, the concern about down-time may outweigh the cost of the installation or change in influencing the decision. The professional computer service industry which carries out and supports installations and upgrades has been rapidly expanding. However, even with such computer professional support, the threat of such down-time coupled with the costs of such professional services is a concern.

[07] As corporate performance and end-user productivity have become increasingly dependent on computers, computer support personnel are continuously under pressure to accomplish more with existing or reduced staff head counts. They are also under pressure to perform tasks as efficiently as possible which may include minimizing effects to existing computer systems and networks.

[08] As shown from the above discussion, there is a need for systems and methods that will improve the ability to manage and support computer systems. Improved systems and methods may enable a person performing computer support to work more efficiently and accomplish more in less time. Benefits may be realized by providing increased functionality to assist in computer maintenance and support.

BRIEF DESCRIPTION OF THE DRAWINGS

[09] The present embodiments will become more fully apparent from the following description and appended claims, taken in conjunction with the accompanying drawings. Understanding that these drawings depict only typical embodiments and are, therefore, not to be considered limiting of the invention's scope, the embodiments will be described with additional specificity and detail through use of the accompanying drawings in which:

[10] Figure 1 is a block diagram of an embodiment of a system for client initiated multicast domain discovery;

[11] Figure 2 is a flow diagram illustrating client initiated multicast domain discovery;

- [12] Figure 3 is a block diagram illustrating data that may be stored and reported to the administrative system;
- [13] Figure 4 is a block diagram of an embodiment of a multicast domain identification;
- [14] Figure 5 is a block diagram of a system being used with multiple subnets;
- [15] Figure 6 is a block diagram illustrating a node connected to multiple networks;
- [16] Figure 7 is a block diagram of a system being used with a Network Address Translation device;
- [17] Figure 8 is a flow diagram of a method for client initiated multicast domain discovery;
- [18] Figure 9 is a flow diagram of an embodiment of a node's method for determining its multicast domain;
- [19] Figure 10 is an exemplary diagram illustrating the alias domains established according to one embodiment of the method for discovering the alias domains; and
- [20] Figure 11 is a block diagram illustrating the major hardware components typically utilized in a computing device.

DETAILED DESCRIPTION

- [21] A method for determining a multicast domain for a node on a computer network is disclosed. A destination is obtained comprising a computer on a computer network. A unique identification is determined for a first router encountered when sending data from a node to the destination. A multicast domain identification is prepared based on the unique identification determined. The multicast domain identification is provided to a central location.
- [22] In an embodiment, a specific port of the first router encountered is determined and a specific port identification is included in the multicast domain identification. The multicast domain identification may comprise a router or similar network device identification and a port identification.
- [23] In a further embodiment the central location may include an administrative system that includes a database of multicast domains.
- [24] In certain embodiments a multicast domain cache may be used. Files may be reported that are in a multicast domain cache.

[25] A computer-readable medium for storing program data is also disclosed. The program data includes executable instructions for implementing a method in a computing device for determining a multicast domain for a node on a computer network. In the method a destination is obtained comprising a computer on a computer network. A unique identification is determined for a first router encountered when sending data from a node to the destination. A multicast domain identification is prepared based on the unique identification determined. The multicast domain identification is provided to a central location.

[26] A system for determining a multicast domain for a node on a computer network is also disclosed. The system includes a computing device and a computer-readable medium. The computer-readable medium is part of the computing device or is in electronic communication with the computing device. The computer-readable medium comprises executable instructions for implementing a method for determining the multicast domain for the computing device. A destination is obtained that comprises a computer on the computer network. A unique identification is determined for a first router encountered when sending data from the computing device to the destination. A multicast domain identification is prepared based on the unique identification determined. The multicast domain identification is sent to an administrative system on the computer network.

[27] It will be readily understood that the components of the embodiments as generally described and illustrated in the Figures herein could be arranged and designed in a wide variety of different configurations. Thus, the following more detailed description of the embodiments of the systems and methods of the present invention, as represented in the Figures, is not intended to limit the scope of the invention, as claimed, but is merely representative of the embodiments of the invention.

[28] The word "exemplary" is used exclusively herein to mean "serving as an example, instance, or illustration." Any embodiment described herein as "exemplary" is not necessarily to be construed as preferred or advantageous over other embodiments. While the various aspects of the embodiments are presented in drawings, the drawings are not necessarily drawn to scale unless specifically indicated.

[29] Several aspects of the embodiments described herein will be illustrated as software modules or components stored in a computing device. As used herein, a software module or component may include any type of computer instruction or computer executable code located within a memory device and/or transmitted as electronic signals over a system bus or network. A software module may, for instance, comprise one or more physical or logical blocks of computer instructions, which may be organized as a routine, program, object, component, data structure, etc., that performs one or more tasks or implements particular abstract data types.

[30] In certain embodiments, a particular software module may comprise disparate instructions stored in different locations of a memory device, which together implement the described functionality of the module. Indeed, a module may comprise a single instruction, or many instructions, and may be distributed over several different code segments, among different programs, and across several memory devices. Some embodiments may be practiced in a distributed computing environment where tasks are performed by a remote processing device linked through a communications network. In a distributed computing environment, software modules may be located in local and/or remote memory storage devices.

[31] Note that the exemplary embodiment is provided as an exemplar throughout this discussion, however, alternate embodiments may incorporate various aspects without departing from the scope of the present invention.

[32] The order of the steps or actions of the methods described in connection with the embodiments disclosed herein may be changed by those skilled in the art without departing from the scope of the present invention. Thus, any order in the Figures or detailed description is for illustrative purposes only and is not meant to imply a required order unless explicitly stated herein.

[33] Although the embodiments herein are described with reference to multicast jobs, they may be applied to broadcast jobs, point to multipoint, and multipoint to multipoint jobs as well.

[34] A multicast alias domain consists of all computers that can see each other's multicast traffic without crossing a router. Another possible definition is that a multicast domain consists of all computers where the IP hop count between any two computers is zero. Because the hop count between all machines in a multicast domain is zero the arrangement of the machines within

the multicast domain does not typically matter. One approach for determining alias domains is described in U.S. Patent Application No. 09/927,708, filed August 10, 2001 and entitled "Method and Apparatus for Dynamically Discovering Alias Domains," hereby incorporated by reference in its entirety. This multicast domain discovery process is controlled from a central server and involves iterating through all subnets to determine the multicast domains.

[35] The embodiments herein use multicast alias domains. However, these embodiments may also use other multipoint communication domains rather than multicast alias domains. For example, the techniques herein may be used with broadcast domains or any other multipoint communication domains.

[36] This system provides a method for multicast domain discovery that allows the individual client computers or nodes to determine their own multicast domain and report the information to a central location, referred to herein as the administrative system 102. Figure 1 is a block diagram of an embodiment of a system 100 for client initiated multicast domain discovery. In an embodiment of the method, the system 100 uniquely identifies the multicast alias domain by the IP address and name of the first router 104 encountered when sending a packet to a known destination, which is the administrative system 102 in Figure 1, from the individual node 106.

[37] A known destination node or computer is used when sending a packet to identify a multicast domain. All of the nodes in the same multicast domain will encounter the same router when sending a packet to the same destination node; thus, they will report the same router name and IP address for the first router encountered and will be labeled as being in the same multicast alias domain.

[38] The multicast alias domain can then be reported to the administrative system 102 through one or more messages, an inventory scan, or similar means. In addition to reporting the alias domain the node 106 can also report files that are currently in the multicast cache, shown in Figure 3.

[39] Figure 2 is a flow diagram 200 illustrating client initiated multicast domain discovery. Some time after the node computer 106 is started 202, it may determine 204 its own multicast alias domain. Once the node 106 has its own multicast alias domain, it may report 206 that information to the administrative system 102. The administrative system 102 may serve as a

central location for storing the information. The information that may be stored and conveyed to the administrative system 102 is shown in FIG. 3.

[40] Figure 3 is a block diagram illustrating data that may be stored and reported to the administrative system 302. The multicast domain identification 308 identifies what multicast domain the node 306 belongs to. The node 306 initially obtains this information and transmits it to an administrative system 302. The administrative system 302 stores the information it receives from the node 306. In most systems there will be many nodes 306 on a computer network reporting their multicast domain identifications 308 to the administrative system 302. As a result, the administrative system 302 typically stores all the multicast domain information in a multicast domain database 310. The multicast domain database 310 may include a plurality of multicast domain records 311 corresponding to the various multicast domains. The multicast domain database 310 may also contain other information not related to multicast domains, such as inventory information. The multicast domain database 310 includes the multicast domain identifications 312 it receives along with node identifications 314, 316, 318 to identify what nodes belong to each multicast domain. Each multicast domain identification 312 may be associated with one or more node identifications 314, 316, 318 for the nodes that belong to that multicast domain.

[41] One or more nodes 306 may use some storage for a multicast domain cache 320. The multicast domain cache 320 comprises storage for electronic files or data that may be needed by the administrative system 302 or other nodes 306. When the node 306 reports its multicast domain identification 308 to the administrative system 302, it may also report the contents of its multicast domain cache 320, if it has one.

[42] Figure 4 is a block diagram of an embodiment of a multicast domain identification 408. The multicast domain identification 408 may be any piece of data that identifies a multicast domain. For example, the multicast domain identification 408 may include a router identification 410 that identifies the first router 104 encountered by a packet sent to the administrative system 102 (or another predetermined destination). The router identification 410 may be a name, number, MAC address, etc. A port identification 412 may identify the port of the router that the packet came in through.

[43] Figure 5 illustrates an embodiment used with multiple subnets 508, 510, 512. In some of the diagrams herein the networks are represented using two octets. The algorithm outlined will work for IP addresses as well. A first subnet 508 and a second subnet 510 are in communication with the router 504 via a first switch 530. A third subnet 512 is in communication with the router 504 via a second switch 532.

[44] The router 504 may have one or more ports 505. The router 504 illustrated has three ports: port 10.0.1.1 505a, port 10.0.2.1 505b and port 10.0.3.1 505c. Even though these ports are identified using an IP address notation, any method of uniquely identifying the ports may be used. Various computers, computing devices and/or network devices may be connected to the router 504.

[45] The nodes 506 in all three subnets 508, 510, 512 will determine the first hop on the route to the administrative system 502. For nodes in the first and second subnets 508, 510 it will be the 10.0.1.1 port 505a of the router 504 and so, in one embodiment, they may report their multicast alias domain using a multicast domain identification as R1-10.0.1.1 (router identification 410 and port identification 412). Nodes 506 in the third subnet 512 will encounter the 10.0.2.1 port 505b of the router 504 when tracing the route to the administrative system 502 and may report their multicast alias domain as R1-10.0.2.1.

[46] The switches 530, 532 do not change the IP hop count, and hence do not impact the multicast domain configuration. Since the nodes 506 are looking for the first network device that decrements the IP address count they will not identify the switch 530, 532. Those familiar with the art will realize that there are known techniques for finding the first network device that decrements the IP address. The Traceroute program demonstrates one such technique for finding the first device that decrements the IP hop count. Details of this technique can be found in TCP/IP Illustrated, Volume 1 by W. Richard Stevens, hereby incorporated by reference in its entirety.

[47] In implementing the algorithm for determining which multicast domain a node belongs to, the systems herein instruct all the nodes to use the same computer (which may be the administrative system) as a destination for determining a route. The first router that is

encountered in the route to the destination is identified and may be used in identifying the multicast domain.

[48] Network environments that contain multiple routers or computers with multiple network cards require all managed nodes to use the same destination when determining their multicast domain. Using the same destination means that they will send their packet to the same well known computer. Without the same destination, issues may arise resulting in unwanted results. Without the same destination, there is no guarantee that the same router will always be encountered. Consider the network shown in Figure 6 which shows a computer 606 with multiple network cards.

[49] In Figure 6, the managed node 606 is connected to three networks 10.0.1.1 608, 10.0.3.1 610 and 10.0.5.1 612. Each of these networks is connected to a router 614, 616, 618. There are three servers. A first server 620 is connected to the first router 614 via network 10.0.2.1 622. A second server 624 is connected to the second router 616 via network 10.0.4.1 626. A third server 628 is connected to the third router 618 via network 10.0.6.1 630. Depending upon the server to which a packet is sent, any one of these three networks may be used. Without a specific well known destination to which a packet is to sent the multicast domain name for the node 606 could be R1-10.0.1.1, R2-10.0.3.1, or R3-10.0.5.1. Once a well known destination, for example the second server 624, is selected only one multicast domain name is possible.

[50] Figure 7 is a block diagram of a system being used with a Network Address Translation device 702. The name of the router may be included in the multicast domain identification 308 to handle the situation where Network Address Translation (NAT) devices 702 are used in the network as shown in the diagram of Figure 7. The NAT device 702 is in electronic communication with a first node 704 and a second node 706 through a first router 708 and a second router 710, as shown. The NAT device 702 makes it possible for two portions of the network to reuse the same network address. Thus as far as the two managed nodes 704, 706 are concerned they are both on network 10.0.1.1 and the address of the router they first encounter is 10.0.1.1. If only the address of the router was used these nodes 704, 706 would incorrectly identify their multicast alias domains. In order to resolve this problem the name of the router is added the multicast alias domain, thus the first node 704 is in the R1-10.0.1.1 multicast alias

domain and the second node 706 is in the R2-10.0.1.1 multicast alias domain, which allows for unique multicast domain names even within a NAT environment.

[51] The name of the router is used for illustration purposes only; any unique attribute of the router could be included in the domain name. Unique attributes of a router may include the name, serial number, MAC address, and so on.

[52] The systems and methods herein of client initiated multicast domain discovery offer the advantage of not needing a centralized server to control the discovery. Additionally because this solution can store the multicast domain information directly in a database there is no need for the administrative system to perform any additional discovery tasks, which greatly reduces the time it takes to perform multicast.

[53] Figure 8 is a method 800 for the system where managed nodes determine their multicast domains. A destination for the nodes to use in determining the multicast domains is provided 802. The individual nodes then determine 804 their own multicast domains and report 806 that information to an administrative system 102.

[54] Figure 9 is a flow diagram of an embodiment of a node's method 900 for determining its multicast domain. The destination for the node to use in determining its multicast domain is obtained 902. The destination may be determined by a system administrator and sent out to all nodes on the network. Alternatively, the nodes 106 may be configured to automatically use the administrative system 102 as the destination. In addition, a volunteer process may be used wherein the first node 106 that is started volunteers to be the destination. Those skilled in the art will appreciate that there are a number of ways in which a destination may be determined.

[55] Then a unique identification for the first router encountered when sending a packet to the destination is determined 904. Those skilled in the art will appreciate that there are a number of ways in which a unique identification for the first router encountered may be determined. For example, the name of the router can be determined by making a `gethostbyaddr` call which makes it possible to determine a hostname given an IP address. Alternately the Address Resolution Protocol (ARP) can be used to determine the MAC address given an IP address.

[56] The port of the router used is also determined 906 in this embodiment. By port the embodiments are attempting to identify which of the routers ports are being used for

communication. This is commonly identified by knowing the IP address of the router encountered. As outlined above, the Traceroute program provides a method to determine the IP address of the first network device that decrements the IP hop count. Then the unique identification and port for the router are provided 908 to a database or an administrative system 102.

[57] Figure 10 is an exemplary block diagram illustrating a computer network on which the present systems and methods may be implemented. In Figure 10, an administrative system 1001 connects to a router 1003 that has three ports: port A 1002, port B 1004 and port C 1006. The administrative system 1001 may be any computer or computing device that has been configured to, or is being used for, storing the aggregated multicast domain information.

[58] The router 1003 may be connected to three switches: a first switch 1005, a second switch 1007 and a third switch 1009. Each switch 1005, 1007, 1009 connects to three subnets. The first switch 1005 connects to three subnets 1010, 1012, and 1014. The second switch 1007 connects to three subnets 1016, 1018, and 1020. The third switch 1009 connects to three subnets 1022, 1024, and 1026. The network nodes or elements 1011, 1013, 1015, 1017, 1019, 1021, 1023, 1025 and 1027 represent computer systems or devices on the computer network. One or more of the nodes may use embodiments of the systems and methods herein for client initiated multicast domain discovery.

[59] Figure 10 is an exemplary diagram illustrating the alias domains established according to one embodiment of the method for discovering the alias domains. Once the nodes have all determined their multicast domains, the administrative system 1001 may store the multicast domain information. A multicast alias domain RPortC 1031 includes subnet A 1010, subnet B 1012, and subnet C 1014. A multicast domain RPortA 1033 includes the subnet D 1016, the subnet E 1018, and the subnet F 1020. A multicast domain RPortB 1035 includes the subnet G 1022, the subnet H 1024, and the subnet I 1026.

[60] Having this information about the network will provide the administrator with detailed information about the network and provide for more intelligent replication of files. For example, using these embodiments it will be possible to determine the most efficient way to cascade a file through the system.

[61] Figure 11 is a block diagram illustrating the major hardware components typically utilized in a computer or computing device used with embodiments herein. Computers and/or computing devices 1102 are known in the art and are commercially available. The major hardware components typically utilized in a computing device 1102 are illustrated in Figure 11. A computing device 1102 typically includes a processor 1103 in electronic communication with input components or devices 1104 and/or output components or devices 1106. The processor 1103 is operably connected to input 1104 and/or output devices 1106 capable of electronic communication with the processor 1103, or, in other words, to devices capable of input and/or output in the form of an electrical signal. Embodiments of devices 1102 may include the inputs 1104, outputs 1106 and the processor 1103 within the same physical structure or in separate housings or structures.

[62] The electronic device 1102 may also include memory 1108. The memory 1108 may be a separate component from the processor 1103, or it may be on-board memory 1108 included in the same part as the processor 1103. For example, microcontrollers often include a certain amount of on-board memory. The memory 1108 may be embodied in RAM, a hard drive, a CD-ROM drive, a DVD-ROM drive, network storage, etc. The memory 1108 is broadly defined as any electronic component capable of storing electronic information.

[63] The processor 1103 is also in electronic communication with a communication interface 1110. The communication interface 1110 may be used for communications with other devices 1102. Thus, the communication interfaces 1110 of the various devices 1102 may be designed to communicate with each other to send signals or messages between the computing devices 1102. A network card or wireless card may be used to implement a communication interface 1110.

[64] The computing device 1102 may also include other communication ports 1112. In addition, other components 1114 may also be included in the electronic device 1102.

[65] Of course, those skilled in the art will appreciate the many kinds of different devices that may be used with embodiments herein. The computing device 1102 may be a one-chip computer, such as a microcontroller, a one-board type of computer, such as a controller, a typical desktop computer, such as an IBM-PC compatible, a Personal Digital Assistant (PDA), a Unix-based workstation, a router, etc. Accordingly, the block diagram of Figure 11 is only meant to

illustrate typical components of a computer or computing device 1102 and is not meant to limit the scope of embodiments disclosed herein.

[66] FIG. 11 illustrates a computer system 1102 that includes or is in electronic communication with a machine-readable medium (memory) on which is stored a set of instructions according to embodiments herein. Although described in the context of a computer system 1102, the embodiments herein may be implemented in any suitable computer system comprising any suitable one or more integrated circuits.

[67] Accordingly, computer system 1102 includes or is in communication with a computer-readable medium on which is stored a set of instructions (i.e., software) embodying any one, or all, of the methodologies described above. For example, software can reside, completely or at least partially, within main memory and/or within processors 1103. For the purposes of this specification, the term “computer-readable medium” shall be taken to include any mechanism that provides (i.e., stores and/or transmits) information in a form readable by a machine (e.g., a computer). For example, a computer-readable medium includes read only memory (ROM); random access memory (RAM); magnetic disk storage media; optical storage media; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc.

[68] Those of skill in the art would understand that information and signals may be represented using any of a variety of different technologies and techniques. For example, data, instructions, commands, information, signals, bits, symbols, and chips that may be referenced throughout the above description may be represented by voltages, currents, electromagnetic waves, magnetic fields or particles, optical fields or particles, or any combination thereof.

[69] Those of skill would further appreciate that the various illustrative logical blocks, modules, circuits, and algorithm steps described in connection with the embodiments disclosed herein may be implemented as electronic hardware, computer software, or combinations of both. To clearly illustrate this interchangeability of hardware and software, various illustrative components, blocks, modules, circuits, and steps have been described above generally in terms of their functionality. Whether such functionality is implemented as hardware or software depends upon the particular application and design constraints imposed on the overall system. Skilled

artisans may implement the described functionality in varying ways for each particular application, but such implementation decisions should not be interpreted as causing a departure from the scope of the present invention.

[70] The various illustrative logical blocks, modules, and circuits described in connection with the embodiments disclosed herein may be implemented or performed with a general purpose processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array signal (FPGA) or other programmable logic device, discrete gate or transistor logic, discrete hardware components, or any combination thereof designed to perform the functions described herein. A general purpose processor may be a microprocessor, but in the alternative, the processor may be any conventional processor, controller, microcontroller, or state machine. A processor may also be implemented as a combination of computing devices, e.g., a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DSP core, or any other such configuration.

[71] The steps of a method or algorithm described in connection with the embodiments disclosed herein may be embodied directly in hardware, in a software module executed by a processor, or in a combination of the two. A software module may reside in RAM memory, flash memory, ROM memory, EPROM memory, EEPROM memory, registers, hard disk, a removable disk, a CD-ROM, or any other form of storage medium known in the art. An exemplary storage medium is coupled to the processor such the processor can read information from, and write information to, the storage medium. In the alternative, the storage medium may be integral to the processor. The processor and the storage medium may reside in an ASIC. The ASIC may reside in a user terminal. In the alternative, the processor and the storage medium may reside as discrete components in a user terminal.

[72] The methods disclosed herein comprise one or more steps or actions for achieving the described method. The method steps and/or actions may be interchanged with one another without departing from the scope of the present invention. In other words, unless a specific order of steps or actions is required for proper operation of the embodiment, the order and/or use of

specific steps and/or actions may be modified without departing from the scope of the present invention.

[73] While specific embodiments and applications of the present invention have been illustrated and described, it is to be understood that the invention is not limited to the precise configuration and components disclosed herein. Various modifications, changes, and variations which will be apparent to those skilled in the art may be made in the arrangement, operation, and details of the methods and systems of the present invention disclosed herein without departing from the spirit and scope of the invention.

[74] What is claimed is: